Pursuing Precision Cosmological Constraints on the Blue Gene/P Supercomputer Platform

J. P. Bernstein,[1] S. E. Kuhlmann,[1] B. Norris,[2] P. M. Ricker[3]

[1]High Energy Physics Division, Argonne National Laboratory, Lemont, IL 60439, USA

[2]Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439, USA

[3]Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA

E-mail: `jpbernst@anl.gov`

Abstract. We present interim results of a two-pronged study of the application of the Argonne National Laboratory Blue Gene/P (BG/P) supercomputer to the study of precision cosmology: simulations of the large-scale structure of the nniverse (computational cosmology) and radiative transfer calculations for Type Ia supernova (SNIa) explosions. We show that porting of the `FLASH` computational cosmology module to the BG/P is complex and that pursuit of simulation runs involving more than 512^3 particles is a significant challenge. We further show that `SEDONA` SNIa radiative transfer calculations scale very well on the BG/P for a minimally parallel case.

1. Introduction

The evidence for the accelerated expansion of the universe [1, 2] represents one of the greatest mysteries of modern science. Several independent observational methods provide evidence pointing to the conclusion that 95% of the current energy density of the universe exists in the form of the dark sector that is different from normal (i.e., baryonic) matter. In particular, the dark sector is believed to be comprised of 22% dark matter, that interacts with baryonic matter ostensibly only through gravity, and 73% dark energy, that experiences gravity as a repulsive force and is believed to be driving cosmic acceleration. An alternative to dark energy that has been suggested is the breakdown of general relativity (a.k.a. modified gravity). In the coming decade, at least four observational probes will be used to improve our understanding of the dark sector: Type Ia supernovae (SNIa), baryon acoustic oscillations (BAO), galaxy cluster counts, and weak gravitational lensing. The age of precision cosmology has arrived, and dark energy constraints on the order of 1% are demanded. While SNIa continue to be a leading probe of dark energy, an improvement in the detailed understanding of their origins is necessary if 1% precision is to be reached.

It is an exciting time to be involved in cosmology because planned astronomical surveys will effectively result in the probes discussed above becoming systematics-limited, making numerical simulations crucial to placing precision constraints on the dark sector. In addition, particle physics theorists predict that the Large Hadron Collider and direct dark matter experiments may soon discover a dark matter candidate particle. If this occurs, a new suite of cosmology simulations will be required to check for consistency with current data and predict indirect dark matter signals. The Dark Energy Survey (DES), a new astronomical survey scheduled to begin science operations in early 2012, will provide the next generation of dark energy and indirect dark matter data. The focus of our research is the production and interpretation of computational cosmology simulations in order to elucidate the nature of the dark sector. SNIa and BAO test the expansion history of the universe and its geometry. Galaxy clusters and weak gravitational lensing are additionally sensitive to the rate at which density fluctuations grow because of gravity and so offer the best hope of distinguishing between modified gravity and dark energy models. Our research addresses theoretical and computational challenges posed by SNIa, BAO, galaxy cluster formation, and weak gravitational lensing. Current and future precision predictions and analyses from these probes must rely on numerical simulations and the systematic error budget will very soon be dominated by theoretical uncertainties. Reducing these uncertainties depends critically on producing more comprehensive and detailed numerical simulations.

We have also engaged in SNIa physics studies. SNIa arise from the thermonuclear detonation of the degenerate remnants of low-mass stars. SNIa are "standard candles," i.e., objects with luminosity deducible from observational properties allowing distances to be calculated. This makes SNIa excellent cosmological tools. observations of distant SNIa in the late 1990s provided the first evidence for the acceleration of cosmic expansion in that these SNe were found to be fainter than predicted by matter-dominated models of the universe. The recent SN data from surveys covering cosmologically relevant distances in the universe, such as the ESSENCE Supernova Survey [3], Supernova Legacy Survey [4], Sloan Digital Sky Survey-II Supernova Survey [5], Supernova Cosmology Project [2], Carnegie Supernova Project [6], and Hubble Higher-z Supernova Search [7], in combination with measurements of the cosmic microwave background anisotropy and BAO, have confirmed accelerated expansion in terms of the dark energy density, Ω_{DE}, and the equation of state parameter, w.

2. Computational Approach

2.1. Computational Cosmology

The accuracy of observations of the large scale structure of the universe from the DES have the potential to significantly distinguish dark energy and modified gravity effects. However, we can only improve on the results of current structure measurements with the help of cosmological simulations that go beyond the current state of the art in terms of physical fidelity and dynamical range. Therefore, we are pursuing the development of a new suite of high-resolution simulations following the growth of cosmic structures through gravitational instability in different cosmological models. These simulations are to include the effects of baryons on structure formation and will have a dynamical range sufficient to replicate the statistics of the forthcoming surveys while simultaneously capturing the internal structure of galaxies and clusters. These simulations are extremely taxing on computational power and memory and will require resources at a level appropriate to a DOE INCITE award. Our computational cosmology studies require hydrodynamic+N-body simulations with a minimum of a 1 Gpc (1 pc $\approx 3.09\times10^{16}$ m) box size, 2048^3 hydrodynamic cells, and 10^{11} dark matter particles. These requirements push the limits of existing codes. We are pursuing science runs with the `FLASH` cosmology module in order to produce the required simulations. Our effort includes the improvement of FLASH cosmology performance on Argonne's Blue Gene/P (BG/P) Intrepid, a 40-rack production machine, and Surveyor, a single test and development rack, with Fusion, an Argonne Linux cluster, used for comparison studies.

2.2. Type Ia Supernovae

The outstanding SN results to date were made possible by the Phillips relation that identified a correlation between SNIa light curve width and luminosity. This allowed the scatter in SNIa intrinsic brightness estimates to be significantly reduced. The theoretical approach to interpreting observational SNIa data involves two main stages: the explosion phase (hydrodynamics + nuclear synthesis), and the free expansion phase involving radiative transfer (RT). In the past decade, much effort has been put into the development of highly sophisticated hydrodynamical models, with which an Argonne–University of Chicago collaboration has successfully reached a world leading position. RT modeling has also been pursued. Work has been performed with codes, such as `SEDONA` [8] (the RT code evaluated in this work), which are capable of processing SNIa remnants in detail. However, the computational demands for multidimensional (multi-D) models have limited the amount of detail that can be taken into account, for example, departures from local thermodynamic equilibrium (LTE). This introduces uncertainties in observational interpretation and obstructs the development of precision dark energy constraints.

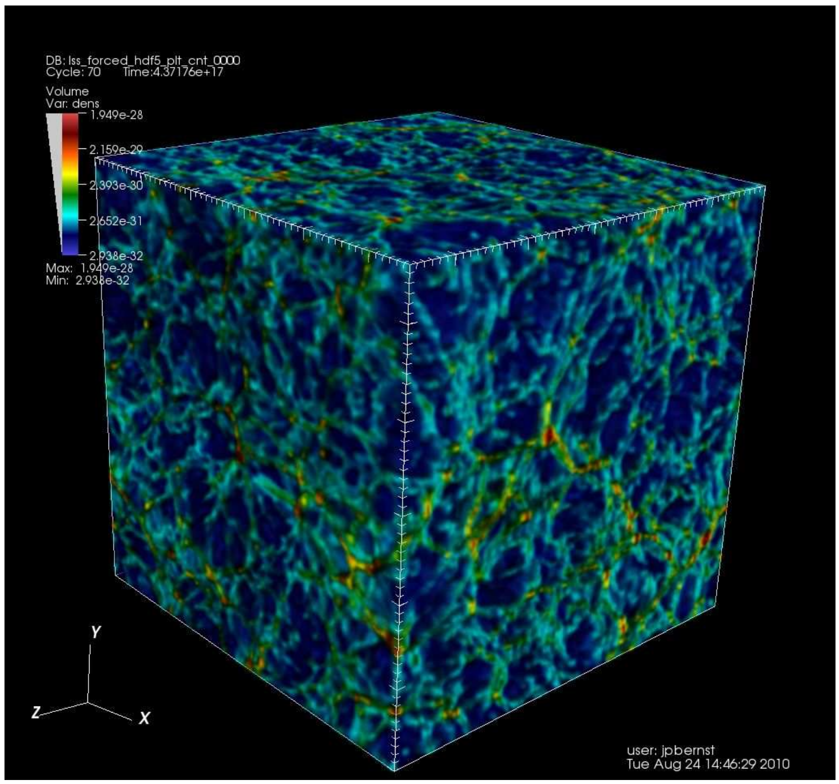

Figure 1. `FLASH` cosmology simulation run at Argonne National Laboratory including 512^3 dark matter particles and hydrodynamics with adaptive mesh refinement. This is believed to be the largest computational cosmology simulation run at Argonne.

3. Computational Results

3.1. Computational Cosmology

`FLASH` is one of the main consumers of Intrepid in the context of SNIa explosion simulations. However, our work is the first sustained effort to utilize Intrepid for `FLASH` computational cosmology simulations. Our experience with `FLASH`, and previously with another computational cosmology code, has shown that porting of such codes from other architectures to BG/P is a complex task. A year-long effort with `FLASH` on Surveyor and Intrepid revealed problems with reading of initial conditions that caused initialization failures for runs larger than 512^3 particles because of inefficient particle loading onto CPUs. With the initialization issue corrected, 1024^3 particle runs were successfully initialized but failed to complete because of memory allocation problems. To date, 512^3 is the largest run that has been successfully completed (see Fig. 1).

However, adaptive-mesh FLASH simulation runs from 128^3 to 512^3 particles have showed a lack of scaling. We traced the cause of the poor scaling to the adaptive mesh refinement implementation. Other adaptive mesh packages are being implemented in the production release of FLASH. In the meantime, we have pursued nonadaptive (unigrid) simulations with the FLASH parallel fast Fourier transform (PFFT) solver and have obtained encouraging scaling results (see Fig. 2).

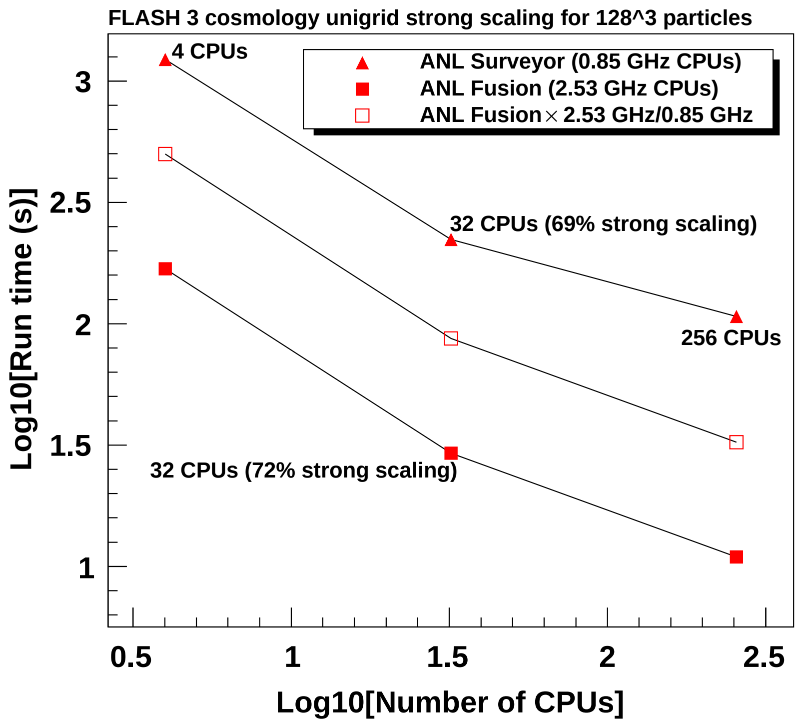

Figure 2. FLASH cosmology strong scaling for 128^3 dark matter particles in unigrid mode with the PFFT solver on Surveyor (VN mode) and Fusion. The scaling from 4 to 32 CPUs is believed to be the closest to FLASH cosmology linear BG/P scaling achieved. The reduction in scaling at 256 CPUs is thought to arise because of communication binding. The open squares show a naive calculation of the Surveyor run time, given the Fusion run time, assuming that the slowdown is due to relative CPU clock speed only.

3.2. Type Ia Supernovae

We showed that SEDONA scales up to 131K CPUs on Intrepid, albeit for a minimally parallel case (see Fig. 3). Progressing the understanding of SNIa systematics calls for large scale simulations requiring an increase in realism via a multi-D RT treatment that moves away from the assumption of LTE. The computational resources required, for example, on the order of a few terabytes of memory, necessitate the utilization of high-performance machines. With upcoming computer architectures, the limits discussed above could be removed. However, multi-D RT simulations are not easily pursued on machines that are compute-optimized. In the non-LTE RT treatment, the physical states at different spatial points depend on each other, which significantly complicates parallelization. In any case, non-LTE calculations are a major source of computational expense. Thus, the creation of an efficient, parallel non-LTE solver is a necessary first step in producing a multi-D RT code that is optimized for the BG/P and can be scaled to future leadership architectures.

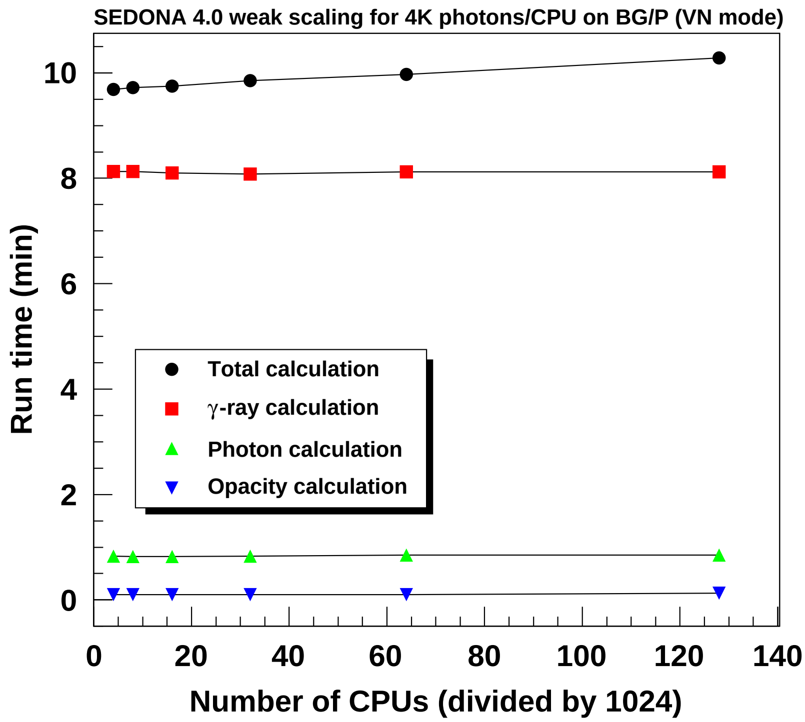

Figure 3. `SEDONA` weak scaling on Intrepid in VN mode (4 out of 4 CPUs per nodes used) for a minimally parallel case.

4. Summary and Conclusion

We have presented a two-pronged approach to addressing the need for increased precision in the investigation of the nature of dark energy and cosmic acceleration: computational cosmology and Type Ia supernova studies. We have contributed to the improvement of the high-performance computing capabilities of both research areas. First, we have undertaken what we believe to be the first sustained effort to utilize BG/P architecture for `FLASH` computational cosmology simulations and have reported on test simulations and scaling trials that revealed significant challenges meriting additional study. Second, we have pursued BG/P tests of the `SEDONA` supernova radiative transfer code and have shown excellent scaling up to the full Intrepid machine for a minimally parallel case.

We thank `SEDONA` PI Dan Kasen for providing access to `SEDONA` source code and problem set-up directions and the University of Chicago `FLASH` Center for code advice. This work was supported in the form of Argonne Computational Postdoctoral Fellowship funds provided by the 2009 United States American Recovery and Reinvestment act.

References

[1] Riess, A. G. et al., 1998, Astronom. J., 116, 1009
[2] Perlmutter, S. et al. 1999, Astrophys. J., 517, 565
[3] Miknaitis, G. et al. 2007, Astrophys. J, 666, 674
[4] Astier, P. et al. 2006, Astronom.. Astrophys., 447, 31
[5] Frieman, J. A. et al. 2008b, Astronom. J., 135, 338
[6] Hamuy, M. et al. 2006, Pub. Astronom. Soc. Pac., 118, 2
[7] Strolger, L.-G. et al. 2004, Astrophys. J., 613, 200
[8] Kasen, D.; Thomas, R. C.; Nugent, P. 2006, Astrophys. J., 651, 366